

WE CLAIM:

- Sub
A1
1. A system for maintaining consistent cached copies of memory in a multi-node computing system having a main memory, comprising:
- 5 at least one memory directory having memory directory entries mapping the main memory, one or more of the memory directory entries including state information for a corresponding line of main memory;
- at least one directory cache for storing directory cache lines corresponding to a subset of the memory directory entries; and
- means for using the state information to allocate memory directory entries to the directory cache.
- 10
2. The system of claim 1, wherein the means for using the state information comprises determining sharing behavior of a memory line corresponding to the state information.
3. The system of claim 1, wherein the state information includes at least one bit in each of the memory directory entries, said at least one bit indicating sharing behavior of its corresponding
- 15 memory lines.
4. The system of claim 1, wherein the state information in a memory directory entry includes data indicating which one or more nodes in the multi-node computing system have an associated data cache that contains a copy of a memory line corresponding to such memory directory entry.

5 The system of claim 1, wherein the state information in a memory directory entry includes a single bit indicator that indicates whether or not a memory directory entry has been accessed by multiple nodes in the multi-node computing system.

6. The system of claim 1 wherein the multi-node computing system further includes a
5 plurality of coherence controller subsystems and wherein each of the at least one directory cache is associated with one or more of the plurality of coherence controller subsystems.

7. The system of claim 1, wherein the subset of the memory directory entries corresponds to a set of most frequently used memory directory entries.

8. A system for maintaining consistent cached copies of memory in a multi-node computing system having a main memory, comprising:

at least one memory directory having memory directory entries mapping the main memory, one or more of the memory directory entries including state information for a corresponding line of main memory;

at least one directory cache for storing directory cache lines corresponding to a subset of
15 the memory directory entries; and

means for using the state information to evict directory cache lines in the directory cache.

9. The system of claim 8, wherein the means for using the state information comprises determining sharing behavior of a memory line corresponding to the state information.

10. The system of claim 8, wherein the state information includes at least one bit in each of the memory directory entries, said at least one bit indicating sharing behavior of its corresponding memory lines.

11. The system of claim 8, wherein the state information in a memory directory entry includes data indicating which one or more nodes in the multi-node computing system have an associated data cache that contains a copy of a memory line corresponding to such memory directory entry.

12. The system of claim 8, wherein the state information in a memory directory entry includes a single bit indicator that indicates whether or not a memory directory entry has been accessed by another node in the multi-node computing system .

13. The system of claim 8 wherein the multi-node computing system further includes a plurality of coherence controller subsystems and wherein each of the at least one directory cache is associated with one or more of the plurality of coherence controller subsystems.

14. The system of claim 8, wherein the subset of the memory directory entries corresponds to a set of most frequently used memory directory entries.

15. Apparatus for maintaining consistent cached copies of memory lines in a multi-node computing system having a main memory, comprising:

at least one memory directory having memory directory entries mapping the main memory; and

memory directory entries, stored in the memory directory, including state information indicating whether one or more nodes in the multi-node computing system have an associated data cache that contains a copy of a memory line corresponding to the memory directory entry.

16. The apparatus of claim 15, further comprising:

at least one directory cache for storing directory cache lines corresponding to a subset of the memory directory entries;

at least one coherence controller, each of the at least one directory cache being associated with one or more of the coherence controllers; and

means for using the state information to allocate directory cache lines to the directory cache.

17. The apparatus of claim 16, wherein the means for using the state information comprises determining sharing behavior of a memory line associated with the state information.

18. The system of claim 16 wherein the state information further includes at least one bit in each of the memory directory entries, said at least one bit indicating sharing behavior of a corresponding memory line.

19. A method of maintaining state information in a memory directory entry of a multi-node computing system having a main memory, one or more nodes in the system having a coherence controller, a processor cache, a memory directory of the main memory and a directory cache of the memory directory, the method comprising the steps of:

receiving a signal at the directory cache in one node of the system indicative of a coherence request for a cached memory line from one of the other nodes of the system; performing a memory directory lookup to determine the location of the directory entry of the cached memory line; and
5 storing information describing sharing behavior of the cached memory line .

20. A method of directory cache eviction in a multi-node computing system having a main memory, one or more of the nodes having a coherence controller, a processor cache, a memory directory associated with the main memory, and a directory cache of the memory directory having directory entries storing state information for associated lines of main memory, the method comprising the step of:

selecting a directory cache line for eviction based on the state information of the retrieved directory cache line; and

updating a validity indicator associated with the selected cache line to indicate that the directory cache line does not include valid information.

15 21. The method of claim 20 wherein the selecting step comprises the steps of:

examining the state information for one or more directory entries in a plurality of directory cache lines and determining sharing behavior of corresponding memory lines; and

selecting for eviction a directory cache line that meets a predetermined sharing behavior criterion.

Sub
121

22. A method of directory cache allocation in a multi-node computing system having a main memory, one or more of the nodes in the system having a coherence controller associated with the main memory, memory directory associated with the main memory, a directory cache of the memory directory having directory entries storing state information for associated lines of main memory, the method comprising the steps of:

retrieving a memory directory entry from a memory directory;

deciding whether to allocate the memory directory entry based on the state information of the memory directory entry.

23. The method of claim 22, wherein the decision to allocate is based on sharing behavior information contained in the memory directory entry.

24. A program storage device, readable by a machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of claim 19.

25. A program storage device, readable by a machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of claim 20.

26. A program storage device, readable by a machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of claim 22.